

AUS9-2000-0483-US1

**CLAIMS**

What is claimed is:

1. A method of balancing a workload across a plurality of servers, the method comprising the steps of:

5 responsive to a request from a requesting client for a distributed service, forwarding the request to a first distributed service manager associated with the requesting client;

10 determining whether the first distributed service manager has information about the distributed service;

15 if the first distributed service manager has information about the distributed service, retrieving the information about the distributed service;

20 if the first distributed service manager does not have information about the distributed service, retrieving information about the distributed service from a second distributed service manager and caching the retrieved information at the first distributed service manager; and

25 sending the retrieved information to the requesting client.

2. The method of claim 1 wherein the first distributed service manager has information about at least two sources for the distributed service and selects a source which will provide best service to the requesting client based on network performance metrics.

AUS9-2000-0483-US1

3. A method of balancing demand for networked services in a distributed data processing system, the method comprising the steps of:

- 5 initializing one or more local service managers within the distributed data processing system, wherein each local service manager provides access to networked services for clients within the distributed data processing system, and wherein each client is uniquely associated with a local service manager;
- 10 initializing one or more distributed service managers within the distributed data processing system, wherein each distributed service manager provides access to networked services to local service managers within the distributed data processing system, and wherein each local service manager is uniquely associated with a distributed service manager;
- 15 receiving, at a distributed service manager, a request for a networked service from a local service manager;
- 20 determining whether the distributed service manager has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service; and
- 25 returning information about a matched networked service from the distributed service manager to the local service manager.

AUS9-2000-0483-US1

4. The method of claim 3 further comprising:  
sending a request for a networked service from a  
requesting client to a local service manager associated  
5 with the requesting client; and

returning information about a matching networked  
service from the local service manager to the requesting  
client, wherein the matching networked service has  
characteristics that match parameters in the request for  
10 a networked service.

5. The method of claim 3 further comprising:  
receiving a request for a networked service at a  
local service manager; and  
15 determining whether the local service manager has  
information about a networked service with  
characteristics that match parameters in the request for  
a networked service.

20 6. The method of claim 5 further comprising:  
if the local service manager has information about a  
matching networked service, returning the information  
about the matching networked service to the requesting  
client;  
25 if the local service manager does not have  
information about a matching networked service,  
forwarding the request for a networked service from the  
local service manager to a distributed service manager  
associated with the local service manager.

AUS9-2000-0483-US1

7. The method of claim 3 further comprising:  
if the distributed service manager has information  
about a matching networked service, returning the  
information about the matching networked service to the  
5 local service manager;  
if the distributed service manager does not have  
information about a matching networked service,  
broadcasting the request for a networked service from the  
distributed service manager to all distributed service  
10 managers in the distributed data processing system;  
receiving information about one or more matching  
networked services at the distributed service manager in  
response to the broadcast request; and  
caching the received information about one or more  
15 matching networked services at the distributed service  
manager.

8. The method of claim 3 further comprising:  
in response to a determination that the distributed  
20 service manager has information about two or more  
matching networked services, selecting a single networked  
service at the distributed service manager.

9. The method of claim 8 further comprising:  
25 performing a load balancing operation at the  
distributed service manager to select the single  
networked service.

10. The method of claim 9 further comprising:  
30 comparing network-related metrics during the load  
balancing operation.

AUS9-2000-0483-US1

11. The method of claim 10 further comprising:  
comparing one or more of network-related metrics  
associated with a network path between a requesting  
client and a providing server.

12. The method of claim 11 wherein the network-related  
metrics are selected from a group comprising:  
bottleneck-link speed, round-trip time, and hop count.

10 13. An apparatus for balancing a workload across a  
plurality of servers, the apparatus comprising:

15 forwarding means for forwarding, responsive to a  
request from a requesting client for a distributed  
service, the request to a first distributed service  
manager associated with the requesting client;

20 determining means for determining whether the first  
distributed service manager has information about the  
distributed service;

25 first retrieving means for retrieving, if the first  
distributed service manager has information about the  
distributed service, the information about the  
distributed service;

second retrieving means for retrieving, if the first  
25 distributed service manager does not have information  
about the distributed service, information about the  
distributed service from a second distributed service  
manager;

30 caching means for caching retrieved information at  
the first distributed service manager; and

AUS9-2000-0483-US1

sending means for sending the retrieved information to the requesting client.

14. The apparatus of claim 13 further comprising:

selecting means for selecting a source which will provide best service to the requesting client based on network performance metrics when the first distributed service manager has information about at least two sources for the distributed service.

10

15. An apparatus for balancing demand for networked services in a distributed data processing system, the apparatus comprising:

first initializing means for initializing one or more local service managers within the distributed data processing system, wherein each local service manager provides access to networked services for clients within the distributed data processing system, and wherein each client is uniquely associated with a local service manager;

second initializing means for initializing one or more distributed service managers within the distributed data processing system, wherein each distributed service manager provides access to networked services to local service managers within the distributed data processing system, and wherein each local service manager is uniquely associated with a distributed service manager;

AUS9-2000-0483-US1

first receiving means for receiving, at a distributed service manager, a request for a networked service from a local service manager;

5 first determining means for determining whether the distributed service manager has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service; and

10 first returning means for returning information about a matched networked service from the distributed service manager to the local service manager.

16. The apparatus of claim 15 further comprising:

15 sending means for sending a request for a networked service from a requesting client to a local service manager associated with the requesting client; and

20 second returning means for returning information about a matching networked service from the local service manager to the requesting client, wherein the matching networked service has characteristics that match parameters in the request for a networked service.

17. The apparatus of claim 15 further comprising:

25 first receiving means for receiving a request for a networked service at a local service manager; and

second determining means for determining whether the local service manager has information about a networked service with characteristics that match parameters in the request for a networked service.

AUS9-2000-0483-US1

18. The apparatus of claim 17 further comprising:  
third returning means for returning, if the local  
service manager has information about a matching  
networked service, the information about the matching  
networked service to the requesting client.

5  
forwarding means for forwarding, if the local  
service manager does not have information about a  
matching networked service, the request for a networked  
service from the local service manager to a distributed  
10 service manager associated with the local service  
manager.

19. The apparatus of claim 15 further comprising:  
fourth returning means for returning, if the  
15 distributed service manager has information about a  
matching networked service, the information about the  
matching networked service to the local service manager;  
broadcasting means for broadcasting, if the  
distributed service manager does not have information  
20 about a matching networked service, the request for a  
networked service from the distributed service manager to  
all distributed service managers in the distributed data  
processing system;  
second receiving means for receiving information  
25 about one or more matching networked services at the  
distributed service manager in response to the broadcast  
request; and  
caching means for caching the received information  
about one or more matching networked services at the  
30 distributed service manager.

AUS9-2000-0483-US1

20. The apparatus of claim 15 further comprising:  
selecting means for selecting, in response to a  
determination that the distributed service manager has  
information about two or more matching networked  
services, a single networked service at the distributed  
service manager.

5  
21. The apparatus of claim 20 further comprising:  
10 performing means for performing a load balancing  
operation at the distributed service manager to select  
the single networked service.

15  
22. The apparatus of claim 21 further comprising:  
first comparing means for comparing network-related  
metrics during the load balancing operation.

20  
23. The apparatus of claim 22 further comprising:  
second comparing means for comparing one or more of  
network-related metrics associated with a network path  
between a requesting client and a providing server.

25  
24. The apparatus of claim 23 wherein the  
network-related metrics are selected from a group  
comprising: bottleneck-link speed, round-trip time, and  
hop count.

30  
25. A computer program product on a computer readable  
medium for use in a data processing system for balancing  
a workload across a plurality of servers, the computer  
program product comprising:

AUS9-2000-0483-US1

forwarding means for forwarding, responsive to a request from a requesting client for a distributed service, the request to a first distributed service manager associated with the requesting client;

5 determining means for determining whether the first distributed service manager has information about the distributed service;

first retrieving means for retrieving, if the first distributed service manager has information about the 10 distributed service, the information about the distributed service;

second retrieving means for retrieving, if the first distributed service manager does not have information about the distributed service, information about the distributed service from a second distributed service manager;

caching means for caching retrieved information at the first distributed service manager; and

sending means for sending the retrieved information 20 to the requesting client.

26. The computer program product of claim 25 further comprising:

25 selecting means for selecting a source which will provide best service to the requesting client based on network performance metrics when the first distributed service manager has information about at least two sources for the distributed service.

AUS9-2000-0483-US1

27. A computer program product on a computer readable medium for use in a data processing system for balancing demand for networked services in a distributed data processing system, the computer program product comprising:

instructions for initializing one or more local service managers within the distributed data processing system, wherein each local service manager provides access to networked services for clients within the distributed data processing system, and wherein each client is uniquely associated with a local service manager;

instructions for initializing one or more distributed service managers within the distributed data processing system, wherein each distributed service manager provides access to networked services to local service managers within the distributed data processing system, and wherein each local service manager is uniquely associated with a distributed service manager;

instructions for receiving, at a distributed service manager, a request for a networked service from a local service manager;

instructions for determining whether the distributed service manager has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service; and

instructions for returning information about a matched networked service from the distributed service manager to the local service manager.

PCT/US2000/0483

AUS9-2000-0483-US1

28. The computer program product of claim 27 further comprising:

instructions for sending a request for a networked service from a requesting client to a local service manager associated with the requesting client; and

instructions for returning information about a matching networked service from the local service manager to the requesting client, wherein the matching networked service has characteristics that match parameters in the request for a networked service.

29. The computer program product of claim 27 further comprising:

instructions for receiving a request for a networked service at a local service manager; and

instructions for determining whether the local service manager has information about a networked service with characteristics that match parameters in the request for a networked service.

30. The computer program product of claim 29 further comprising:

instructions for returning, if the local service manager has information about a matching networked service, the information about the matching networked service to the requesting client;

instructions for forwarding, if the local service manager does not have information about a matching networked service, the request for a networked service

AUS9-2000-0483-US1

from the local service manager to a distributed service manager associated with the local service manager.

31. The computer program product of claim 27 further

5 comprising:

instructions for returning, if the distributed service manager has information about a matching networked service, the information about the matching networked service to the local service manager;

10 instructions for broadcasting, if the distributed service manager does not have information about a matching networked service, the request for a networked service from the distributed service manager to all distributed service managers in the distributed data processing system;

15 instructions for receiving information about one or more matching networked services at the distributed service manager in response to the broadcast request; and

20 instructions for caching the received information about one or more matching networked services at the distributed service manager.

32. The computer program product of claim 27 further

comprising:

25 instructions for selecting, in response to a determination that the distributed service manager has information about two or more matching networked services, a single networked service at the distributed service manager.

AUS9-2000-0483-US1

33. The computer program product of claim 32 further comprising:

instructions for performing a load balancing operation at the distributed service manager to select  
5 the single networked service.

34. The computer program product of claim 33 further comprising:

instructions for comparing network-related metrics  
10 during the load balancing operation.

35. The computer program product of claim 34 further comprising:

instructions for comparing one or more of  
15 network-related metrics associated with a network path between a requesting client and a providing server.

36. The computer program product of claim 35 wherein the network-related metrics are selected from a group comprising: bottleneck-link speed, round-trip time, and  
20 hop count.

DRAFT ATTACHMENT